

UNITED STATES PATENT APPLICATION

FOR

HYBRID TIME DIVISION MULTIPLEXING AND
DATA TRANSPORT

INVENTORS:

DAVID CHEN
FRANK MARRONE
MATS FRANNHAGEN

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN, LLP
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1026

(503) 684-6200

EXPRESS MAIL NO. EL625195445US

HYBRID TIME DIVISION MULTIPLEXING AND DATA TRANSPORT

FIELD OF THE INVENTION

[0001] The invention relates to network switches. More particularly, the invention relates to transport of both time division multiplexed (TDM) traffic and network protocol (e.g., ATM, IP) traffic.

BACKGROUND OF THE INVENTION

[0002] In high bandwidth networks such as fiber optic networks, lower bandwidth services such as voice communications are aggregated and carried over a single fiber optic link. However, because the aggregated data can have different destinations some mechanism for switching the aggregated components is required. Switching can be performed at different levels of aggregation.

[0003] Current switching is accomplished in a synchronous manner. Signals are routed to a cross-connect or similar switching device that switch and route signals at some predetermined granularity level, for example, byte by byte. Synchronous switching in a cross-connect is a logically straight forward method for switching. However, because data flow between network nodes is not necessarily consistent, switching bandwidth may not be used optimally in a synchronous cross-connect. One source of data may use all available bandwidth while a second source of data may transmit data sporadically.

[0004] In order to support data sources that transmit at or near peak bandwidth, cross-connects are designed to provide the peak bandwidth to all data sources because specific data rates of specific data sources are not known when the cross-connect is designed. As

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements.

Figure 1 is a block diagram of one embodiment of a line card for use in transporting TDM traffic and data traffic.

Figure 2 is a conceptual block diagram of one embodiment of a time slot manager and associated components.

Figure 3 is a block diagram of a dumb card accessing an ATM/IP engine on an intelligent card.

Figure 4 illustrates one embodiment of a set of line interface cards interconnected via a backplane to provide a time slot interchange.

INS
C1

DETAILED DESCRIPTION

[0006] Methods and apparatuses for transporting of both time division multiplexed (TDM) traffic and network protocol traffic are described. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the invention. It will be apparent, however, to one skilled in the art that the invention can be practiced without these specific details. In other instances, structures and devices are shown in block diagram form in order to avoid obscuring the invention.

[0007] Reference in the specification to "one embodiment" or "an embodiment" means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the invention. The appearances of the phrase "in one embodiment" in various places in the specification are not necessarily all referring to the same embodiment.

[0008] In one embodiment, a network switch includes multiple interface cards and a backplane that interconnects the interface cards. The interface cards receive network traffic and perform time and line switching on the data. The network traffic can include a combination of time division multiplexed (TDM) data and network data (e.g., ATM cells, IP packets). In one embodiment, the channels that carry network traffic to the interface cards are pre-configured as either TDM channels or network channels. The channels are processed as appropriate for their respective types by the interface cards. Because both TDM and network traffic can be processed by a single interface card, the number of cards within the network switch can be reduced.

[0009] **Figure 1** is a block diagram of one embodiment of a line card for use in transporting TDM traffic and data traffic. The line card of Figure 1 is coupled between a telecommunications line (e.g., optical fiber) and a backplane or switching fabric of a network switch. For reasons of simplicity neither the telecommunications line nor the backplane are illustrated in Figure 1. The telecommunications line can be any type of telecommunications line known in the art. The backplane can be a type of backplane known in the art or a backplane specifically designed to operate with the line cards described herein.

[0010] Line interface 100 is coupled to telecommunications or other network lines. Line interface 100 can be, for example, an optical fiber interface that includes optical to electrical conversion circuitry to generate electrical signals that represent optical data received via the optical fiber(s). Line interface 100 can be coupled to one or more lines. Line interface 100 can also be coupled to receive electrical signals.

[0011] Physical layer framer 110 receives the electrical signals output by line interface 100 and generates frames of data. Physical layer framer 110 organizes the data received through line interface 100 into frames having predetermined formats so that the data can be processed by the components of the card as well as other cards of the network switch (not shown in Figure 1).

[0012] For example, physical layer framer 110 can convert data between Synchronous Optical Network (SONET) frames and an internal cell format that is used by the components of the card of Figure 1. The SONET standard is described in the American National Standards Institute (ANSI) standards T1.105 and T1.106 and in the Bellcore Technical Recommendations TR-TSY-000253. Other conversions can also be supported.

sub
a1

[0013] Time Slot Manager (TSM) 120 receives and transmits frames of data from and to physical layer framer 110 and routes the data between physical layer framer 110 the appropriate serializer/deserializer (SERDES). TSM 120 also schedules the transmission of data based on, for example, data type, available bandwidth and/or other considerations. Data scheduling is described in greater detail in U.S. Patent application number 09/XXX,XXX (Atty. Docket No. P013), filed _____, and entitled "DISTRIBUTED CONTROL OF DATA FLOW IN A NETWORK SWITCH," which is assigned to the corporate assignee of the present U.S. Patent application and incorporated by reference herein.

[0014] TSM 120 sends data to one of a group serializer/deserializers for switching of data. TSM 120 can send data to multiple serializer/deserializers to multicast or broadcast data to multiple sources. In most situations, serializer/deserializers (e.g., 130, 132, 134) transmit data across a backplane, or switching fabric, to another card (not shown in Figure 1) within the network switch. Alternatively, as described in greater detail below a serializer/deserializer can loop the data back to the card from which the data is sent. The serializer/deserializers also receive data from the backplane.

[0015] Asynchronous Transfer Mode/Internet Protocol (ATM/IP) engine 160 is coupled to TSM 120 and includes cell/packet switching engine 140 and ATM/POS framer 150. ATM/IP engine 160 is used for processing of ATM or IP data only. TDM data, for example, is scheduled and processed by TSM 120. In one embodiment, ATM/IP engine 160 resides on a daughter card that can be coupled with an interface to TSM 120. By having ATM/IP engine 160 on a daughter card, basic cards can be manufactured and ATM/IP functionality can be added to the card by coupling the ATM/IP engine daughter

card to the main card. Alternatively, ATM/IP engine 160 can also be an integrated component of the line interface card of Figure 1.

[0016] Cell/packet switching engine 140 provides switching at the cell/packet level. Data received from TSM 120 is switched as necessary and sent to ATM/POS framer 150 for framing in the appropriate format. The framed data is sent from ATM/POS framer 150 to TSM 120 for routing to the appropriate serializer/deserializer.

[0017] **Figure 2** is a conceptual block diagram of one embodiment of a time slot manager and associated components. TSM 120 of Figure 1 is illustrated as two TSM blocks, one for ingress processing (220) and one for egress processing (225). In one embodiment, data flow through the components of Figure 2 can be configured in a per STS-1 channel basis. As described in greater detail below, ingress and/or egress flow can include 48 or more STS-1 channels that can be either TDM traffic or ATM/IP traffic.

[0018] In one embodiment, ingress TSM 220 includes 20 ingress Time Slot Interchanges (TSIs); however any number of TSIs can be provided. The ingress TSIs (e.g., 230, 232, 234) are coupled to receive data from physical layer framer 110 (e.g., SONET framed data, ATM cells, IP packets). The ingress TSIs are also coupled to receive data from POS/ATM framer 150. In one embodiment, the ingress TSIs receive STS-1 formatted channels from one or more telecommunications lines via physical layer framer 110 and STS-1 formatted channels can be looped back via egress TSM 225 through POS/ATM framer 150.

[0019] The ingress TSIs perform time switching on the incoming channels. Time switching involves switching the order in which data is transmitted or the times at which the data is transmitted. The time switched channels output by the ingress TSIs are input

to cell multiplexer (cell MUX) 240. In one embodiment, the ingress TSIs communicate data via fixed length cells whether or not the protocol used for the data uses fixed length cells. In other words, IP packets are communicated within the TSM and on the backplane using fixed length cells.

[0020] In one embodiment, cell multiplexer 240 receives fixed length cells from the ingress TSIs and demultiplexes the cells to the appropriate serializer/deserializer so that the cells can be communicated via a backplane link (not shown in Figure 2). In one embodiment, the demultiplexing is based on cell header information for the respective cells. The cell header information can include, for example, a destination card identifier, an interface identifier, a channel identifier, etc.

[0021] The cells output by cell multiplexer 240 are input to serializer/deserializers (e.g., SERDES 250, SERDES 252, SERDES 254), which are coupled to the respective outputs of cell multiplexer 240. The serializer/deserializers convert the incoming cells to serial streams of data to be transported via a backplane.

[0022] On the egress side, cell demultiplexer (CELL DEMUX) 245 receives cell and TDM data from serializer/deserializers (e.g., SERDES 260, SERDES 262, SERDES 264) coupled to the backplane. The serializer/deserializers receive data from corresponding serializer/deserializers on another card that is coupled to the backplane. The serializer/deserializers convert serial data to cell data.

[0023] In one embodiment, line switching egress TSI 250 receives up to 1056 STS-1 channels from cell demultiplexer 245. Of the 1056 STS-1 channels, up to 96 are selected by line switching egress TSI 250. SONET line switching is provided by line switching egress TSI 250.

[0024] Path switching egress TSI 255 receives up to 96 STS-1 channels from line switching egress TSI 250 and outputs up to 48 STS-1 channels. In one embodiment, path switching egress TSI 250 presents up to 48 STS-1 channels to physical layer framer 110 and up to 48 STS-1 channels to POS/ATM framer 150. STS-1 level path switching is provided by path switching egress TSI 255. The channels sent from path switching egress TSI 255 to physical layer framer 110 are transmitted over telecommunications lines coupled to line interface 110.

[0025] With respect to ingress TDM traffic flows from an external device to the backplane, which is left to right in Figure 2, Egress traffic flows from the backplane to the external device, which is right to left in Figure 2. Line interface 100 and physical layer framer 110 provide STS-1 channels carrying TDM data to ingress TSM 220. Ingress TSM 220 routes the channels to the appropriate serializer/deserializer for transmission across the backplane to a target card. Traffic that is destined for the same card (e.g., to a different STS-1 channel to a different physical interface on the same card) are looped back by ingress TSM 220 to egress TSM 225 via the backplane.

[0026] STS-1 channels from the backplane are processed by egress TSM 225 for line switching as well as path switching. Up to 48 STS-1 channels are selected and sent to physical layer framer 110 for transmission to line interface 100.

[0027] The components of Figure 2 can also be used for processing cell/packet data flow. In the ingress direction, STS-1 channels from line interface 100 and physical layer framer 110 and cells/packets are passed through ingress TSM 220 to the SERDES link destined for the same card. In other words, the incoming cells/packets are looped back via the backplane.

[0028] Egress TSM 225 sends the cells/packets looped back via the backplane to POS/ATM framer 150. In an ATM application, ATM cells are carried within a 57-byte backplane cell. A 4-byte header is added to carry information for scheduling.

[0029] In a POS application, IP packets are switched across the backplane by variable-length cells. In one embodiment, the variable length cells can be from 40 to 127 bytes. The backplane super cells from the ingress switching carry the destination card and port information in the header, which is used by cell demultiplexer 240 in ingress TSM 220 to route the cells to the various SERDES links.

[0030] The looped back cells/packets are sent from POS/ATM framer 150 to cell MUX 240 through cell/packet switching engine 140. Cell MUX 240 routes the cell/packets to the appropriate SERDES for forwarding to a target card across the backplane.

[0031] In the egress direction, cells or packets are received from the backplane by SERDES components and sent to cell demultiplexer 245, which routes the cells to egress line switching egress TSI 250 and to the ATM/IP engine for egress switching and processing. After egress switching and processing, ATM cells or IP packets, based on the header information that identifies the egress logical port (or STS-1 channel), are inserted into the appropriate STS-1 channel by the POS/ATM framer.

[0032] In one embodiment, the STS-1 channels from the POS/ATM framer are looped through the SERDES link destined to the same card to the egress TSM 225 before being sent to the physical layer framer 110 and line interface 100 for transmission. In one embodiment, for every cell or IP packet transmitted, a feedback signal is sent to the

source interface card for scheduling purposes. For example, the feedback information can be carried in a cell header.

[0033] **Figure 3** is a block diagram of a dumb card accessing an ATM/IP engine on an intelligent card. In one embodiment, traffic carried on a dumb interface card can be processed by an ATM or IP layer on another interface card having an ATM/IP engine. Unused STS-1 channels on the ATM/POS framer of the intelligent card are used for channels from the dumb card.

[0034] For example, on an OC-48 card coupled to a 48 UPSR or BLSR ring, only 24 STS-1 channels are used. The unused 24 STS-1 channels can be used to process channels from a dumb interface card. A channel of data (e.g., STS-1, STS-3, STS-12) is received by line interface 300 of dumb interface card 375. The channel is processed by physical layer framer 310 and TSM 320 as described above. TSM 320 routes the data to one of one of the serializer/deserializers (SERDES 330, 332, 334) of dumb interface card 375, which transmits the data over backplane 390 to one of the serializer/deserializers (SERDES 130, 132, 134) on intelligent interface card 350.

[0035] The data received by intelligent interface card 350 from dumb interface card is routed by TSM 120 to ATM/IP engine 160 for processing in the manner described above. TSM 120 receives the processed data from ATM/IP engine 160 and routes the processed data to the appropriate serializer/deserializer for transmission over backplane 390 to dumb interface card 375.

[0036] The processed data received by a serializer/deserializer of dumb interface card 325 from intelligent interface card 350 is looped back through TSM 320 to one of the

serializer/deserializers of dumb interface card 375. The data is then transmitted via backplane 390 to a target interface card (not shown in Figure 3).

[0037] **Figure 4** illustrates one embodiment of a set of line interface cards interconnected via a backplane to provide a time slot interchange. The example of Figure 4 is a 960 x 960 STS-1 level time slot interchange; however, any size interchange can be provided in a similar manner.

[0038] In one embodiment, the switches of figure 1 include multiple cards that are interconnected by a switching fabric. In one embodiment, the cards have both an ingress data path and an egress data path. The ingress data path is used to receive data from a network and transmit the data to an appropriate card within the switch. The egress data path is used to receive data from the switching fabric and transmit the data across the network.

[0039] Each ingress interface card (e.g., 400, 405, 410) includes an ingress TSI (e.g., 450, 455, 460) that receives data input channels from an external source. In one embodiment, the data is TDM data; however, data can be in any format, for example, IP packets or ATM cells. The ingress TSIs are coupled to ingress serializer/deserializers. In one embodiment, each ingress card has a serializer/deserializer for each egress card to which the ingress card is coupled.

[0040] Each egress interface card (e.g., 415, 420, 425) includes a serializer/deserializer to for each ingress interface card to which the egress interface card is coupled. The serializer/deserializer of the egress interface card are coupled to an egress TSI (e.g., 465, 470, 475). The egress serializer/deserializers are coupled to an egress TSI that outputs data to a device external to the egress card.

sub
ar

[0041] Because each ingress card is coupled to each egress card, the interconnection between the ingress cards and the egress cards has n^2 connections where n is the number of ingress/egress cards. Thus, the interconnection is referred to as an " n^2 mesh," or an " n^2 switching fabric." The mesh is described in greater detail in U.S. Patent application number _____, entitled "A FULL MESH INTERCONNECT BACKPLANE ARCHITECTURE," filed December 22, 2000, which is assigned to the corporate assignee of the present application and incorporated by reference.

[0042] In one embodiment, each backplane link between an ingress interface card and an egress interface card can carry up to 48 STS-1 channels. As mentioned above, in one embodiment, each interface card includes 20 ingress TSIs. Thus, the bandwidth provided by a 20 TSI interface card is 960 STS-1 channels. By changing the number of TSIs and the number of interconnections across the backplane, the number of STS-1 channels supported can be modified. A protocol for use in communicating over the mesh is described in greater detail in U.S. Patent application number (P005) _____, entitled "A BACKPLANE PROTOCOL," filed December 22, 2000, which is assigned to the corporate assignee of the present invention and incorporated by reference.

[0043] In one embodiment, in the ingress direction, the ingress TSI of each ingress interface card routes the 48 channels received by the ingress interface card to the appropriate serializer/deserializer. The serializer/deserializers transmit data across the backplane to the appropriate egress interface card. In the egress direction each egress interface card receives 1056 channels (960 from the backplane and 96 from loopback). The egress TSI selects and routes 48 of the channels to an external device.

[0044] In the following example of routing channels with a distributed TSI, the system consists of two OC-3 interface cards supporting a total of 6 STS-1 channels. The three STS-1 channels of interface card 1 are channel 1, channel 2 and channel 3. The three STS-1 channels of interface card 2 are channel 4, channel 5 and channel 6. The TSI functions to be implemented are:

Ingress STS-1 channel	Egress STS-1 channel
1	5
2	6
3	1
4	2
5	4
6	3

Table 1: TSI function

[0045] To fulfill the TSI function of Table 1, the ingress TSI and egress TSI tables for cards 1 and 2 are configured as described in Tables 2 and 3, and Tables 4 and 5, respectively.

Ingress STS channel	Backplane/loopback channel
1	5
2	6
3	1
--	--
--	--
--	--

Table 2: Ingress TSI Table for Card 1

Backplane/loopback channel	Egress STS channel
1	--
2	1
3	3
4	2
5	--
6	--

Table 3: Egress TSI Table for Card 1

Ingress STS channel	Backplane/loopback channel
--	--
--	--
--	--
4	2
5	4
6	3

Table 4: Ingress TSI Table for Card 2

Backplane/loopback channel	Egress STS channel
1	6
2	--
3	--
4	--
5	4
6	5

Table 5: Egress TSI Table for Card 2

[0046] Thus, each ingress TSI and egress TSI has an associated table or tables that indicate the routing of channels between the ingress side and the egress side. Other types of data, for example, ATM or IP data can be routed between the ingress TSIs and egress TSIs in a similar manner except that the data is routed through the ATM/IP engine as described above.

[0047] In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes can be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.